CLAIMS:

1. A host-fabric adapter installed at a host system for connecting to a switched fabric of a data network, comprising:

a Micro-Engine (ME) arranged to establish connections and support data transfers via said switched fabric;

a serial interface arranged to receive and transmit data packets from said switched fabric for data transfers; and

a host interface arranged to receive and transmit host data transfer requests, in the form of descriptors, from said host system for data transfers, and incorporated therein a host interface Hardware Assist (HWA) mechanism configured to pre-process host descriptors for descriptor format errors in parallel with descriptor fetches so as to offload said Micro-Engine (ME) from having to check for said descriptor format errors.

2. The host-fabric adapter as claimed in claim 1, wherein said host interface Hardware Assist (HWA) mechanism comprises:

a Descriptor Format Checker arranged to check host descriptors from said host system for said descriptor format errors using predetermined descriptor format rules and descriptor contents in response to a ME instruction from said Micro-Engine (ME); and

a Descriptor Register Array arranged in parallel with said Descriptor Format Checker to

supply descriptor status information to said Micro-Engine (ME) when descriptor fetching operations are completed.

3. The host-fabric adapter as claimed in claim 2, wherein said host descriptors from said host system provide information needed to complete send/receive, remote direct memory access (RDMA) write/read operations for data transfers, and include send/receive descriptors utilized to control transmission/reception of a single data packet, and remote direct memory access (RDMA) descriptors utilized to additionally indicate the address of remote information.

4. The host-fabric adapter as claimed in claim 1, further comprising:

a Receive FIFO interface arranged to receive data packets from said switched fabric via said serial interface; and

a Transmit FIFO interface arranged to transmit data packets to said switched fabric via said serial interface.

5. The host-fabric adapter as claimed in claim 4, wherein said Receive FIFO interface incorporates therein a Protection Index and Offset Hardware Assist (HWA) mechanism configured to process the Virtual Address (VA) and Memory Handle (MH) of an incoming data packet, via the serial interface, and generate therefrom a Protection Index (PI) and Offset so as to offload said Micro-Engine (ME) from processing data packets for RDMA read/write operations.

6. The host-fabric adapter as claimed in claim 5, wherein said Protection Index and Offset Hardware Assist (HWA) mechanism comprises:

a packet Buffer arranged to temporarily store an incoming data packet from the serial interface;

a packet Loading Logic arranged to start loading the data packet from said packet Buffer in response to an ME instruction from said Micro-Engine (ME);

a Protection Index and Offset Logic arranged to calculate the Protection Index (PI) and Offset based on the Virtual Address (VA) and associated Memory Handle (MH) of the incoming data packet from the packet Buffer in accordance with a load Virtual Address (VA) request from said packet Loading Logic; and

a Multiplexer arranged to select, as an output, the Protection Index (PI) and Offset needed for each task of said Micro-Engine (ME) in response to said ME instruction from said Micro-Engine (ME).

7. The host-fabric adapter as claimed in claim 6, wherein said packet Buffer is a single packet first-in/first-out (FIFO) storage device.

8. The host-fabric adapter as claimed in claim 6, wherein said Protection Index (PI) and Offset are obtained by said Protection Index and Offset Logic using the following formula:

Offset = VA (11:0); and
Protection Index (PI) = VA (43:12) - MH (31:0),

where the Offset is the lower 12-bits of the Virtual Address (VA) of the incoming data packet to indicate which bytes within a single page are being addressed.

9. The host-fabric adapter as claimed in claim 6, wherein said Protection Index and Offset Logic comprises:

a plurality of task Virtual Address Registers arranged to receive the Virtual Address (VA) from the data packet and the load Virtual Address request in response to an ME instruction from said Micro-Engine (ME);

task Multiplexers arranged to obtain the Offset from the Virtual Address (VA) from the data packet previously registered in different Virtual Address Registers;

a Subtractor arranged to subtract the Memory Handle (MH) of the data packet input from the Virtual Address (VA) of the data packet previously;

a plurality of task Protection Index Registers arranged to load the result of the subtraction in accordance with the load Protection Index (PI) request and the ME instruction from said Micro-Engine (ME); and

an output Multiplexer arranged to select between outputs of different task Protection Index Registers as the Protection Index (PI) in response to the ME instruction from said Micro-Engine (ME).

10. The host-fabric adapter as claimed in claim 6, wherein said Protection Index and Offset Logic comprises:

a first Register arranged to receive the data packet and generate therefrom the Offset in accordance with the load Virtual Address request;

a Subtractor arranged to subtract the Memory Handle (MH) from the Virtual Address (VA) of the data packet; and

a second Register arranged to generate the Protection Index (PI) based on the result of the subtraction in accordance with the load Protection Index (PI) request.

11. The host-fabric adapter as claimed in claim 4, further comprising:

an address translation interface which provides an interface for address translation, and which is addressable by write data and system controls from said Micro-Engine (ME), via a system data bus and a system control bus;

a context memory which provides an interface to a context manager, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data bus and said system control bus, for providing the necessary context for a work queue pair used for sending and receiving data packets;

a local bus interface which provides an interface to a local bus, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data bus and said system control bus, for supporting system accessible context connections and data transfers; and

a completion queue/doorbell manager interface which provides an interface to completion queues, and doorbell and memory registration rules, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data bus and said system control bus.

12. The host-fabric adapter as claimed in claim 4, wherein said Micro-Engine (ME) comprises:

one or more Data Multiplexers arranged to supply appropriate interface data based on an ME instruction;

an Instruction Memory arranged to provide said ME instruction based on downloadable microcode;

an Arithmetic Logic Unit (ALU) arranged to perform mathematical, logical and shifting operations, and supply write data to the host interface, the address translation interface, the context memory interface, the local bus interface, the completion queue/doorbell manager interface, the Receive FIFO interface and the Transmit FIFO interface, via said system write data bus; and

an Instruction Decoder arranged to supply system controls to the host interface, the address translation interface, the context memory interface, the local bus interface, the completion queue/doorbell manager interface, the Receive FIFO interface and the Transmit FIFO interface, via said system control bus, to execute said ME instruction from said Instruction Memory to

control operations of said Data Multiplexers, and to determine functions of said Arithmetic Logic Unit (ALU).

13. The host-fabric adapter as claimed in claim 12, wherein said Instruction Memory corresponds to a static random-access-memory (SRAM) provided to store microcode that are downloadable for providing said ME instruction to said Instruction Decoder.

14. The host-fabric adapter as claimed in claim 1, wherein said host interface, said serial interface and said Micro-Engine (ME) are configured in accordance with the "*Virtual Interface (VI) Architecture Specification*", the *"Next Generation Input/Output (NGIO) Specification*" and the *"InfiniBand™ Specification*".

15. A host-fabric adapter installed at a host system for connecting to a switched fabric of a data network, comprising:

a Micro-Engine (ME) arranged to establish connections and support data transfers via said switched fabric;

a serial interface arranged to receive and transmit data packets from said switched fabric for data transfers;

a host interface arranged to receive and transmit host data transfer requests, in the form of descriptors, from said host system for data transfers; and

a first-in/first-out (FIFO) interface arranged to receive and transmit data packets to/from said switched fabric via said serial interface, and incorporated therein a Protection Index and Offset Hardware Assist (HWA) mechanism configured to process the Virtual Address (VA) and Memory Handle (MH) of data packets and generate therefrom a Protection Index (PI) and Offset so as to offload said Micro-Engine (ME) from having to process all data packets for data transfers.

16. The host-fabric adapter as claimed in claim 15, wherein said Protection Index and Offset Hardware Assist (HWA) mechanism comprises:

a packet Buffer arranged to temporarily store an incoming data packet from the serial interface;

a packet Loading Logic arranged to start loading the data packet from said packet Buffer in response to an ME instruction from said Micro-Engine (ME);

a Protection Index and Offset Logic arranged to calculate the Protection Index (PI) and Offset based on the Virtual Address (VA) and associated Memory Handle (MH) of the incoming data packet from the packet Buffer in accordance with a load Virtual Address (VA) request from said packet Loading Logic; and

a Multiplexer arranged to select, as an output, the Protection Index (PI) and Offset needed for each task of said Micro-Engine (ME) in response to said ME instruction from said Micro-Engine (ME).

17. The host-fabric adapter as claimed in claim 16, wherein said packet Buffer is a single packet first-in/first-out (FIFO) storage device.

18. The host-fabric adapter as claimed in claim 16, wherein said Protection Index (PI) and Offset are obtained by said Protection Index and Offset Logic using the following formula:

Offset = VA (11:0); and
Protection Index (PI) = VA (43:12) - MH (31:0),

where the Offset is the lower 12-bits of the Virtual Address (VA) of the incoming data packet to indicate which bytes within a single page are being addressed.

19. The host-fabric adapter as claimed in claim 6, wherein said Protection Index and Offset Logic comprises:

a plurality of task Virtual Address Registers arranged to receive the Virtual Address (VA) from the data packet and the load Virtual Address request in response to an ME instruction from said Micro-Engine (ME);

task Multiplexers arranged to obtain the Offset from the Virtual Address (VA) from the data packet previously registered in different Virtual Address Registers;

a Subtractor arranged to subtract the Memory Handle (MH) of the data packet input from the Virtual Address (VA) of the data packet previously;

a plurality of task Protection Index Registers arranged to load the result of the subtraction in accordance with the load Protection Index (PI) request and the ME instruction from said Micro-

Engine (ME); and

an output Multiplexer arranged to select between outputs of different task Protection Index Registers as the Protection Index (PI) in response to the ME instruction from said Micro-Engine (ME).

20. The host-fabric adapter as claimed in claim 16, wherein said Protection Index and Offset Logic comprises:

a first Register arranged to receive the data packet and generate therefrom the Offset in accordance with the load Virtual Address request;

a Subtractor arranged to subtract the Memory Handle (MH) from the Virtual Address (VA) of the data packet; and

a second Register arranged to generate the Protection Index (PI) based on the result of the subtraction in accordance with the load Protection Index (PI) request.

21. The host-fabric adapter as claimed in claim 15, wherein said host interface includes a host interface Hardware Assist (HWA) mechanism configured to pre-process host descriptors for descriptor format errors in parallel with descriptor fetches so as to offload said Micro-Engine (ME) from having to check for said descriptor format errors.

22. The host-fabric adapter as claimed in claim 21, wherein said host interface

Hardware Assist (HWA) mechanism comprises:

a Descriptor Format Checker arranged to check host descriptors from said host system for said descriptor format errors using predetermined descriptor format rules and descriptor contents in response to a ME instruction from said Micro-Engine (ME); and

a Descriptor Register Array arranged in parallel with said Descriptor Format Checker to supply descriptor status information to said Micro-Engine (ME) when descriptor fetching operations are completed.

23. The host-fabric adapter as claimed in claim 22, wherein said host descriptors from said host system provide information needed to complete send/receive, remote direct memory access (RDMA) write/read operations for data transfers, and include send/receive descriptors utilized to control transmission/reception of a single data packet, and remote direct memory access (RDMA) descriptors utilized to additionally indicate the address of remote information.

24. The host-fabric adapter as claimed in claim 15, further comprising:

an address translation interface which provides an interface for address translation, and which is addressable by write data and system controls from said Micro-Engine (ME), via a system data bus and a system control bus;

a context memory which provides an interface to a context manager, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data

bus and said system control bus, for providing the necessary context for a work queue pair used for sending and receiving data packets;

a local bus interface which provides an interface to a local bus, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data bus and said system control bus, for supporting system accessible context connections and data transfers; and

a completion queue/doorbell manager interface which provides an interface to completion queues, and doorbell and memory registration rules, and which is addressable by write data and system controls from said Micro-Engine (ME), via said system data bus and said system control bus.

25. The host-fabric adapter as claimed in claim 24, wherein said Micro-Engine (ME) comprises:

one or more Data Multiplexers arranged to supply appropriate interface data based on an ME instruction;

an Instruction Memory arranged to provide said ME instruction based on downloadable microcode;

an Arithmetic Logic Unit (ALU) arranged to perform mathematical, logical and shifting operations, and supply write data to the host interface, the address translation interface, the context memory interface, the local bus interface, the completion queue/doorbell manager interface, the Receive FIFO interface and the Transmit FIFO interface, via said system write data

bus; and

an Instruction Decoder arranged to supply system controls to the host interface, the address translation interface, the context memory interface, the local bus interface, the completion queue/doorbell manager interface, the Receive FIFO interface and the Transmit FIFO interface, via said system control bus, to execute said ME instruction from said Instruction Memory to control operations of said Data Multiplexers, and to determine functions of said Arithmetic Logic Unit (ALU).

26. The host-fabric adapter as claimed in claim 25, wherein said Instruction Memory corresponds to a static random-access-memory (SRAM) provided to store microcode that are downloadable for providing said ME instruction to said Instruction Decoder.

27. The host-fabric adapter as claimed in claim 15, wherein said host interface, said serial interface and said Micro-Engine (ME) are configured in accordance with the "*Virtual Interface (VI) Architecture Specification*", the *"Next Generation Input/Output (NGIO) Specification*" and the *"InfiniBand™ Specification*".

28. A host-fabric adapter, comprising:

a Micro-Engine (ME) arranged to establish connections and support data transfers via a switched fabric;

a host interface arranged to receive and transmit host data transfer requests, in the form of descriptors, from said host system for data transfers, and incorporated therein a host interface Hardware Assist (HWA) mechanism configured to pre-process host descriptors for descriptor format errors in parallel with descriptor fetches so as to offload said Micro-Engine (ME) from exclusively checking for said descriptor format errors;

a serial interface arranged to receive and transmit data packets from said switched fabric for data transfers; and

a first-in/first-out (FIFO) interface arranged to receive and transmit data packets to/from said switched fabric via said serial interface, and incorporated therein a Protection Index and Offset Hardware Assist (HWA) mechanism configured to process the Virtual Address (VA) and Memory Handle (MH) of data packets and generate therefrom a Protection Index (PI) and Offset so as to offload said Micro-Engine (ME) from exclusively processing data packets for data transfers.

29. The host-fabric adapter as claimed in claim 28, wherein said host interface Hardware Assist (HWA) mechanism comprises:

a Descriptor Format Checker arranged to check host descriptors from said host system for said descriptor format errors using predetermined descriptor format rules and descriptor contents in response to a ME instruction from said Micro-Engine (ME); and

a Descriptor Register Array arranged in parallel with said Descriptor Format Checker to

supply descriptor status information to said Micro-Engine (ME) when descriptor fetching operations are completed.

30. The host-fabric adapter as claimed in claim 29, wherein said host descriptors from said host system provide information needed to complete send/receive, remote direct memory access (RDMA) write/read operations for data transfers, and include send/receive descriptors utilized to control transmission/reception of a single data packet, and remote direct memory access (RDMA) descriptors utilized to additionally indicate the address of remote information.

31. The host-fabric adapter as claimed in claim 28, wherein said Protection Index and Offset Hardware Assist (HWA) mechanism comprises:

a packet Buffer arranged to temporarily store an incoming data packet from the serial interface;

a packet Loading Logic arranged to start loading the data packet from said packet Buffer in response to an ME instruction from said Micro-Engine (ME);

a Protection Index and Offset Logic arranged to calculate the Protection Index (PI) and Offset based on the Virtual Address (VA) and associated Memory Handle (MH) of the incoming data packet from the packet Buffer in accordance with a load Virtual Address (VA) request from said packet Loading Logic; and

a Multiplexer arranged to select, as an output, the Protection Index (PI) and Offset needed

for each task of said Micro-Engine (ME) in response to said ME instruction from said Micro-Engine (ME).

32. The host-fabric adapter as claimed in claim 31, wherein said packet Buffer is a single packet first-in/first-out (FIFO) storage device.

33. The host-fabric adapter as claimed in claim 31, wherein said Protection Index (PI) and Offset are obtained by said Protection Index and Offset Logic using the following formula:

Offset = VA (11:0); and
Protection Index (PI) = VA (43:12) - MH (31:0),

where the Offset is the lower 12-bits of the Virtual Address (VA) of the incoming data packet to indicate which bytes within a single page are being addressed.

34. The host-fabric adapter as claimed in claim 31, wherein said Protection Index and Offset Logic comprises:

a plurality of task Virtual Address Registers arranged to receive the Virtual Address (VA) from the data packet and the load Virtual Address request in response to an ME instruction from said Micro-Engine (ME);

task Multiplexers arranged to obtain the Offset from the Virtual Address (VA) from the data packet previously registered in different Virtual Address Registers;

a Subtractor arranged to subtract the Memory Handle (MH) of the data packet input from

the Virtual Address (VA) of the data packet previously;

a plurality of task Protection Index Registers arranged to load the result of the subtraction in accordance with the load Protection Index (PI) request and the ME instruction from said Micro-Engine (ME); and

an output Multiplexer arranged to select between outputs of different task Protection Index Registers as the Protection Index (PI) in response to the ME instruction from said Micro-Engine (ME).

35. The host-fabric adapter as claimed in claim 31, wherein said Protection Index and Offset Logic comprises:

a first Register arranged to receive the data packet and generate therefrom the Offset in accordance with the load Virtual Address request;

a Subtractor arranged to subtract the Memory Handle (MH) from the Virtual Address (VA) of the data packet; and

a second Register arranged to generate the Protection Index (PI) based on the result of the subtraction in accordance with the load Protection Index (PI) request.

36. The host-fabric adapter as claimed in claim 1, wherein said host interface, said serial interface, and said FIFO interface and said Micro-Engine (ME) are configured in accordance with the "*Virtual Interface (VI) Architecture Specification*", the *"Next Generation Input/Output*

(NGIO) Specification" and the "*InfiniBand™ Specification*".